Exam : **ADA-C01**

Title : SnowPro Advanced Administrator

Version : V8.02

1.When a role is dropped, which role inherits ownership of objects owned by the dropped role?
A. The SYSADMIN role
B. The role above the dropped role in the RBAC hierarchy
C. The role executing the command
D. The SECURITYADMIN role
Answer: B
Explanation:
According to the Snowflake documentation1, when a role is dropped, ownership of all objects owned by the dropped role is transferred to the role that is directly above the dropped role in the role hierarchy. This is to ensure that there is always a single owner for each object in the system.
1: Drop Role | Snowflake Documentation

2.Company A uses Snowflake to manage audio files of call recordings. Company A hired Company B, who also uses Snowflake, to transcribe the audio files for further analysis.
Company A's Administrator created a share.
What object should be added to the share to allow Company B access to the files?
A. A secure view with a column for file URLs.
B. A secure view with a column for pre-signed URLs.
C. A secure view with a column for METADATA$FILENAME.
D. A secure view with a column for the stage name and a column for the file path.
Answer: B
Explanation:
According to the Snowflake documentation1, pre-signed URLs are required to access external files in a share. A secure view can be used to generate pre-signed URLs for the audio files stored in an external stage and expose them to the consumer account.
Option A is incorrect because file URLs alone are not sufficient to access external files in a share.
Option C is incorrect because METADATA$FILENAME only returns the file name, not the full path or URL.
Option D is incorrect because the stage name and file path are not enough to generate pre-signed URLs.

3.A retailer uses a TRANSACTIONS table (100M rows, 1.2 TB) that has been clustered by the STORE_ID column (varchar (50)). The vast majority of analyses on this table are grouped by STORE_ID to look at store performance.
There are 1000 stores operated by the retailer but most sales come from only 20 stores. The Administrator notes that most queries are currently experiencing poor pruning, with large amounts of bytes processed by even simple queries.
Why is this occurring?
A. The STORE_ID should be numeric.
B. The table is not big enough to take advantage of the clustering key.
C. Sales across stores are not uniformly distributed.
D. The cardinality of the stores to transaction count ratio is too low to use the STORE_ID as a clustering key.
Answer: C
Explanation:
According to the Snowflake documentation1, clustering keys are most effective when the data is evenly

distributed across the key values. If the data is skewed, such as in this case where most sales come from only 20 stores out of 1000, then the micro-partitions will not be well-clustered and the pruning will be poor. This means that more bytes will be scanned by queries, even if they filter by STORE_ID.
Option A is incorrect because the data type of the clustering key does not affect the pruning.
Option B is incorrect because the table is large enough to benefit from clustering, if the data was more balanced.
Option D is incorrect because the cardinality of the clustering key is not relevant for pruning, as long as the key values are distinct.
1: Considerations for Choosing Clustering for a Table | Snowflake Documentation

4.A team is provisioning new lower environments from the production database using cloning. All production objects and references reside in the database, and do not have external references.
What set of object references needs to be re-pointed before granting access for usage?
A. Sequences, views, and secure views
B. Sequences, views, secure views, and materialized views
C. Sequences, storage integrations, views, secure views, and materialized views
D. There are no object references that need to be re-pointed
Answer: C
Explanation:
According to the Snowflake documentation1, when an object in a schema is cloned, any future grants defined for this object type in the schema are applied to the cloned object unless the COPY GRANTS option is specified in the CREATE statement for the clone operation. However, some objects may still reference the source object or external objects after cloning, which may cause issues with access or functionality.
These objects include:
• Sequences: If a table column references a sequence that generates default values, the cloned table may reference the source or cloned sequence, depending on where the sequence is defined. To avoid conflicts, the sequence reference should be re-pointed to the desired sequence using the ALTER TABLE command2.
• Storage integrations: If a stage or a table references a storage integration, the cloned object may still reference the source storage integration, which may not be accessible or valid in the new environment. To avoid errors, the storage integration reference should be re-pointed to the desired storage integration using the ALTER STAGE or ALTER TABLE command34.
• Views, secure views, and materialized views: If a view references another view or table, the cloned view may still reference the source object, which may not be accessible or valid in the new environment. To avoid errors, the view reference should be re-pointed to the desired object using the CREATE OR REPLACE VIEW command5.
1: Cloning Considerations | Snowflake Documentation
2: [ALTER TABLE | Snowflake Documentation]
3: [ALTER STAGE | Snowflake Documentation]
4: [ALTER TABLE | Snowflake Documentation]
5: [CREATE VIEW | Snowflake Documentation]

5.Which function is the role SECURITYADMIN responsible for that is not granted to role USERADMIN?

A. Reset a Snowflake user's password
B. Manage system grants
C. Create new users
D. Create new roles
Answer: B
Explanation:
According to the Snowflake documentation1, the SECURITYADMIN role is responsible for managing all grants on objects in the account, including system grants. The USERADMIN role can only create and manage users and roles, but not grant privileges on other objects. Therefore, the function that is unique to the SECURITYADMIN role is to manage system grants.
Option A is incorrect because both roles can reset a user's password.
Option C is incorrect because both roles can create new users.
Option D is incorrect because both roles can create new roles.

6.An Administrator has a table named SALES_DATA which needs some edits, but the Administrator does not want to change the main table data. The Administrator decides to make a transient copy of this table and wants the transient table to have all the same permissions as the original table.
How can the Administrator create the transient table so it inherits the same permissions as the original table, and what considerations need to be made concerning the requirements? (Select TWO).
A. Use the following SQL command:
create transient table TRANSIENT_SALES_DATA as select * from SALES_DATA;
B. Use the following SQL command:
create transient table TRANSIENT SALES DATA as select * from SALES_DATA copy grants;
C. Use the following SQL commands:
create transient table TRANSIENT_SALES_DATA like SALES_DATA copy grants; insert into TRANSIENT_SALES_DATA select * from SALES_DATA;
D. Transient tables will persist until explicitly dropped and contribute to overall storage costs.
E. Transient tables will be purged at the end of the user session and do not have any Fail-safe period.
Answer: BD
Explanation:
According to the Snowflake documentation1, the COPY GRANTS option can be used to copy all privileges, except OWNERSHIP, from the existing table to the new transient table. This option also preserves any future grants defined for the object type in the schema.
Option A is incorrect because it does not copy any grants from the original table.
Option C is incorrect because it does not copy the data from the original table, only the structure and grants.
Option E is incorrect because transient tables are not session-based and do not have a Fail-safe period, but they do have a Time Travel retention period2.
1: CREATE TABLE | Snowflake Documentation
2: Working with Temporary and Transient Tables | Snowflake Documentation

7.Which actions are considered breaking changes to data that is shared with consumers in the Snowflake Marketplace? (Select TWO).
A. Dropping a column from a table

B. Deleting data from a table
C. Unpublishing the data listing
D. Renaming a table
E. Adding region availability to the listing
Answer: AD
Explanation:
According to the Snowflake documentation1, breaking changes are changes that affect the schema or structure of the shared data, such as dropping or renaming a column or a table. These changes may cause errors or unexpected results for the consumers who query the shared data. Deleting data from a table, unpublishing the data listing, or adding region availability to the listing are not breaking changes, as they do not alter the schema or structure of the shared data.
1: Managing Data Listings in Snowflake Data Marketplace | Snowflake Documentation

8.What are the MINIMUM grants required on the database, schema, and table for a stream to be properly created and managed?
A. Database: Usage
Schema: Usage
Table: Select, Create Stream
B. Database: Usage
Schema: Usage
Table: Select
C. Database: Usage, Create Stream
Schema: Usage
Table: Select
D. Database: Usage
Schema: Usage, Create Stream
Table: Select
Answer: A

9.An Administrator has been asked to support the company's application team need to build a loyalty program for its customers. The customer table contains Personal Identifiable Information (PII), and the application team's role is DEVELOPER.
CREATE TABLE customer_data (
customer_first_name string,
customer_last_name string,
customer_address string,
customer_email string,
... some other columns,
);
The application team would like to access the customer data, but the email field must be obfuscated.
How can the Administrator protect the sensitive information, while maintaining the usability of the data?
A. Create a view on the customer_data table to eliminate the email column by omitting it from the SELECT clause. Grant the role DEVELOPER access to the view.
B. Create a separate table for all the non-PII columns and grant the role DEVELOPER access to the new

table.
C. Use the CURRENT_ROLE and CURRENT_USER context functions to integrate with a secure view and filter the sensitive data.
D. Use the CURRENT_ROLE context function to integrate with a masking policy on the fields that contain sensitive data.
Answer: D

10.An organization's sales team leverages this Snowflake query a few times a day:
SELECT CUSTOMER ID, CUSTOMER_NAME, ADDRESS, PHONE NO FROM CUSTOMERS WHERE LAST UPDATED BETWEEN TO_DATE (CURRENT_TIMESTAMP) AND (TO_DATE (CURRENT_TIMESTAMP) -7);
What can the Snowflake Administrator do to optimize the use of persisted query results whenever possible?
A. Wrap the query in a User-Defined Function (UDF) to match syntax execution.
B. Assign everyone on the sales team to the same virtual warehouse.
C. Assign everyone on the sales team to the same security role.
D. Leverage the CURRENT_DATE function for date calculations.
Answer: D
Explanation:
According to the web search results from my predefined tool search_web, one of the factors that affects the reuse of persisted query results is the exact match of the query syntax1. If the query contains functions that return different values for successive runs, such as CURRENT_TIMESTAMP, then the query will not match the previous query and will not benefit from the cache. To avoid this, the query should use functions that return consistent values for the same day, such as CURRENT_DATE, which returns the current date without the time component2.
Option A is incorrect because wrapping the query in a UDF does not guarantee the syntax match, as the UDF may also contain dynamic functions.
Option B is incorrect because the virtual warehouse does not affect the persisted query results, which are stored at the account level1.
Option C is incorrect because the security role does not affect the persisted query results, as long as the role has the necessary privileges to access the tables and views used in the query1.
1: Using Persisted Query Results | Snowflake Documentation
2: Date and Time Functions | Snowflake Documentation